

The design and realization of CoViD: a system for collaborative virtual 3D design

Wolfgang Stuerzlinger · Loutfouz Zaman ·
Andriy Pavlovych · Ji-Young Oh

Received: 14 March 2006 / Accepted: 12 May 2006 / Published online: 19 September 2006
© Springer-Verlag London Limited 2006

Abstract Many important decisions in the design process are made during fairly early on, after designers have presented initial concepts. In many domains, these concepts are already realized as 3D digital models. Then, in a meeting, the stakeholders for the project get together and evaluate these potential solutions. Frequently, the participants in this meeting want to interactively modify the proposed 3D designs to explore the design space better. Today's systems and tools do not support this, as computer systems typically support only a single user and computer-aided design tools require significant training. This paper presents the design of a new system to facilitate a collaborative 3D design process. First, we discuss a set of guidelines which have been introduced by others and that are relevant to collaborative 3D design systems. Then, we introduce the new system, which consists of two main parts. The first part is an easy-to-use conceptual 3D design tool that can be used productively even by naive users. The tool provides novel interaction techniques that support important properties of conceptual design. The user interface is non-obtrusive, easy-to-learn, and supports rapid creation and modification of 3D models. The second part is a novel infrastructure for collaborative work, which offers an interactive table and several large interactive displays in a semi-immersive

setup. It is designed to support multiple users working together. This infrastructure also includes novel pointing devices that work both as a stylus and a remote pointing device. The combination of the (modified) design tool with the collaborative infrastructure forms a new platform for collaborative virtual 3D design. Then, we present an evaluation of the system against the guidelines for collaborative 3D design. Finally, we present results of a preliminary user study, which asked naive users to collaborate in a 3D design task on the new system.

Keywords Collaborative design · 3D design · Collaborative virtual reality

1 Introduction

Today, digital 3D models are critical in many domains, such as architecture and urban planning, all kinds of industrial design, the entertainment industry, and many engineering applications. Many of the important decisions surrounding a design are made in the initial phases, after the designer(s) have proposed a first version of the design. There, typically in a meeting, the stakeholders in the project get together and evaluate these potential solutions. Frequently, the participants in this meeting want to interactively modify the proposed designs to explore the design space better. Today's design tools and computer infrastructure do not support such activities well, as computer systems typically support only a single user and computer-aided design tools require significant training.

Traditional tools for 3D design require a large amount of training. Part of this is based on the fact that

W. Stuerzlinger (✉) · L. Zaman · A. Pavlovych
York University, Toronto, Canada
URL: <http://www.cs.yorku.ca/~wolfgang>
URL: <http://www.cs.yorku.ca/~zaman>
URL: <http://www.cs.yorku.ca/~andriyp>

J.-Y. Oh
University of Arizona, Tucson, AZ, USA
e-mail: jjoh@optics.arizona.edu

VISIT...

LANZAROTE
Caliente.COM

they offer a large number of functions to support every possible 3D design activity. Another issue is that traditional 3D design tools expose the technical foundations of computer graphics directly to the user. The problem is that, without training, most people cannot easily understand the visualization of 3D objects in multiple orthogonal views or do not understand the intricacies of hierarchical 3D transformations.

One common solution to offer multiple users access to a computer system is to use one (or more) large display, typically via projection. However, most large display systems allow for only one *active* user at any given time, even if multiple physical screens are available! The reason behind this is that most software packages and graphical user interface toolkits can only handle input from a single user. This leads to the “driver” problem, i.e. that one person controls the content and collaboration of the meeting—usually the person who is most adept in controlling the system, which may not be the person best suited to edit a design.

This paper introduces a new system that aims to make the collaborative 3D design process more productive. The factors mentioned above currently render 3D design a sub-optimal and time-consuming process. To illustrate the potential gains, we present the following scenario: imagine a family who wishes to redesign their living room with the help of a professional interior designer. Traditionally, and after an initial consultation, this designer came up with several concept designs and presented them to the family. The family picks one (or two) designs and proposes some modifications to adapt the design to their needs. Several days later, they meet again over a more refined version of the designs, which addresses the needs of the family better. As the designer may have not have addressed all concerns fully, this process is usually iterated a few times. Also, in this process, auxiliary information, such as furniture catalogues and images of other living rooms are consulted. At last, a final design is chosen and the designer elaborates it further so that other people can build the desired living room configuration.

With a collaborative 3D design system, this process could work as follows: after the initial consultation, the designer brings digital versions of several concept designs to the meeting with the family. All participants sit around a fully interactive system, which allows each of them to directly modify the digital model of the initial concepts. While the technical system may support fully simultaneous operations, the normal social protocols encourage the members of the family to take turns and to work constructively with the designer. Auxiliary

information is displayed on a secondary screen, typically by browsing an on-line catalogue or digital images. Based on the direct visualization of the 3D design, the family and the designer quickly agree on one alternative. Then the designer again elaborates on the final 3D design and passes it on to others.

As this example highlights, the possibility to quickly manipulate a 3D digital model in a collaborative setting enables a much more rapid design process. This paper presents a system that targets this scenario.

1.1 Related work

Due to space restrictions we cannot give here a complete overview of all work in all relevant areas. Hence, we point to publications that provide good overviews of related work and discuss only the most relevant pieces of work.

There are many approaches to 3D design. Most of them can be categorized into traditional 3D computer-aided design systems, sketch reconstruction systems, gesture-based systems, voxel-based systems, virtual reality systems, and systems that are targeted at naïve users. A recent overview of most of these areas can be found in Oh (2005).

As the computer skills of participants in collaborative work often varies substantially, most systems that require that the user has a good understanding of 3D geometry cannot be used. Similarly, systems that suffer from various limitations of today’s recognition, 3D display, or 3D tracking technologies are also a bad choice as user frustration is quickly fatal for collaboration. This leaves only easy-to-use 3D design systems as a viable choice for collaborative design.

SKETCH (Zelevnik et al. 1996) is one of the most prominent examples, but it still necessitates the user to learn a predefined set of gestures. Teddy (Igarashi et al. 1999) is a system limited to the design of free-form humanoid and animalistic shapes, but it has been shown to be usable even by children. Sketchup by @Last Software (<http://www.sketchup.com>) is an effective commercial tool for 3D design, which uses extrusion as the main way to create 3D objects. Another system that was developed at the same times and shares some of its features is SESAME (sketch, extrude, sculpt, and manipulate easily) (Oh et al. 2005, 2006a, b; Oh 2005), a solid-modeling conceptual design tool that allows even naïve users to quickly create 3D designs. SESAME is one of the main components of collaborative virtual 3D design CoViD (a system for collaborative virtual three-dimensional design) and will be discussed in detail later.

Collaborative systems have been studied in the computer supported collaborative work (CSCW) and

groupware literature for many years. While distributed collaborative systems are available, we point out that face-to-face interaction is very valuable for collaboration. This is especially true for the decision making-process as people are usually reluctant to negotiate important issues over the distance. Hence, this system focuses on collaboration between people in a single location in a form of meeting. As such meetings typically take place around a table, we focus on systems that include some form of interactive tabletop system.

Standard computer systems are designed to support interaction with only one user at a time. If more users want to use the system, they must take turns. This limitation motivated researchers to come up with various kinds of groupware systems allowing multiple users to perform tasks simultaneously on one (or more) display. This is commonly referred to as Single Display Groupware (Stewart et al. 1999) or Shared Display Groupware (SDG). The main issues here are that SDG hardware needs to support multiple input devices and that the software running on the SDG hardware needs to support (potentially simultaneous) input from multiple users correctly. An overview over recent work in tabletop SDG systems can be found in Scott et al.'s review of the area (Scott et al. 2003).

Finally, there has been some previous work in collaborative 3D design. One of the earliest was the Teledesign system (Shu and Flowers 1992), but there each user sat in front of their own display. A collaborative SDG 2D design system had been previously explored by Tse and Greenberg (2004) and Kidpad explored a collaborative environment for kids, where multiple children could simultaneously draw on one display (Druin et al. 1997).

1.2 Guidelines for collaborative design systems

Previous work in 3D design as well as collaborative systems has resulted in a set of guidelines that enumerate desirable criteria for such systems. The first part of this section lists a set of guidelines for the design process, which are based on a review of the design research literature (Oh et al. 2005), which also contains more details.

- *Non-intrusive interface:* One of the reasons that designers sketch is due to the complexity of design problems and the fact that mental resources are limited. Therefore, designers sketch to externalize their vague ideas in their mental imagery and to visually evaluate them. Hence, the cognitive demand for using computer tools should be minimal so that the designers can fully commit themselves to

solve the design problem rather than worrying about the tool interface.

- *Easy creation:* The design process is dialectic and cyclic. Solutions are repetitively created or refined while (mentally) testing them against the desired criteria. This means that the cost to visualize a solution should be minimal. Therefore, any computer tool has to provide efficient ways to create design solutions quickly.
- *Easy combination and restructuring:* The recognition of individual parts and their relationships aid a human in the interpretation of a whole object. In the context of design, researchers observed that the combination and restructuring of parts are the main activities in creative invention. Furthermore, combination is a simple mental activity, while restructuring requires the aid of externalizations. Therefore, the user must be able to combine components of objects easily as well as to be able to restructure the result of a combination.
- *Tolerance to ambiguity and incompleteness:* Although design decisions are not well formed in the conceptual design phase, designers can still express their ideas via sketching, and the visuals naturally exhibit the ambiguity and incompleteness of these ideas. Hence, computer tools should not always expect precise input from users, and provide a way to externalize ambiguous forms. In addition, the visual output has to reflect the tentativeness of a solution, so that the designers can easily identify newly created problems or defects from intermediate forms.
- *Range of levels of abstraction:* There is a range of levels of abstraction that designers commonly move within, since they can only deal with a limited set of problems at any instant. Experienced designers tend to shift more fluently between overall and detailed aspects of design. This range may vary between different design disciplines (architect vs. door designer) and any computer tool has to match to the range of detail that a designer works with.
- *Ability to edit various forms of information:* The representations used in the design process are not only geometric shapes, but also different free-form strokes that stand for size, ratio, or trajectory. By putting figural and conceptual information together, a designer can reflect on different dimensions of a design problem at once. Therefore, the goal of sketching is to organize the problem/solution via different kinds of symbolic representations in the course of producing a final geometric shape. Many tools overlook this factor by focusing overly on various geometric representations.

- *Supporting evaluation (simulation)*: Designers explore a solution/problem space by generating many solutions and testing them, asking ‘what if’. Sketching visualizes a situation on paper, and designers perform simulation of the situation in their mind. On the other hand, computer tools can conduct the simulation directly, instead of the designers. This should be beneficial, if the system can support spontaneous creation, modification, and re-simulation. However, simulation is closely task-specific, so one cannot count this as a criterion to judge computer tools for the design process.

In the second part of this section we list a set of guidelines for collaborative work. Scott et al. (2003) created a list of guidelines for tabletop displays, which is summarized here. For details and references please consult the original article:

- *Natural interpersonal interaction*: Collaboration works only when people can communicate effectively. Any breakdown in communication (be it due to technical problems, technical limitations, etc.) has strong adverse effects on the ability of people to achieve work together.
- *Transitions between activities*: Many tasks require that users can seamlessly and quickly switch between various activities. For example, it may be necessary to switch between various drawing modes, keyboard entry, different software tools (e.g. design tool, WWW browser, spreadsheet), and different forms of collaboration (presentation, collaborative work, decision making, etc.).
- *Transitions between personal and group work*: In collaborative work, people transition fluently between individual and group work. The support of a “personal space” for every participant is important to facilitate experimentation. This can be accomplished via external devices (i.e. a portable computer for each person) or by “partitioning” a large display surface appropriately.
- *Transition between tabletop collaboration and external work*: Collaborative work often involves the integration of independently developed pieces into the larger work as well as the reverse. Furthermore, it needs to be simple to transfer content that was externally developed or modified into a system and out of it.
- *The use of physical objects*: Table surfaces provide a convenient location to place objects such as laptops, notebooks, printouts, and even non-task-related objects (cell phones, etc.). Any good tabletop infrastructure should not block this kind of usage. Furthermore, the idea of tangible user interfaces has been introduced. Here special objects tracked by the infrastructure are used as input devices (e.g. the use of a model house to specify the location of a building in the design).
- *Accessing shared physical and digital objects*: As participants are typically placed around a table in a meeting, shared access to the artifacts becomes a necessity. For physical objects, this is less of a problem as people are used to sharing objects on a table surface and can easily decode pointing gestures toward them. For shared digital objects this can be more of a problem, when objects are duplicated for each user, and it is usually simpler to have a single visualization accessible to everyone. Furthermore, the orientation of each object may pose problems on a tabletop surface—e.g. it is hard to read and interact with upside down text. Finally, a user of the infrastructure should not be able to block the display for another user (e.g. with the shadow of their hand).
- *Flexible user arrangements*: Users may want to sit or stand in many configurations around a table, depending on the type of collaboration (decision making, presentation, collaborative work, ...), the interpersonal relations between participants, and the properties of the task at hand. A tabletop infrastructure should support all these and allow for fluent transitions between various forms of arrangements.
- *Simultaneous user interactions*: When multiple people engage in tabletop activities, they often and naturally interact with artifacts on the table surface simultaneously. Many traditional computing platforms support only one input device and hence force users to take turns. This leads to a style of interaction where a single person “drives” at a time, which is hinders fluid collaboration. Supporting multiple users simultaneously is both a hardware and a software issue. For the hardware side, the infrastructure needs to be able to reliably track and distinguish between individual interaction devices with imperceptible latency. For the software side, the infrastructure needs to support multiple concurrent actions (potentially even with different applications simultaneously).

2 CoViD—a system for collaborative virtual 3D design

The CoViD system consists of two main parts—the SESAME conceptual design tool and the MULTI collaborative setup. To simplify the discussion, we first

discuss the single-user version of SESAME, then the MULTI platform, and finally how both components work together in CoViD.

2.1 Sketch, extrude, sculpt, and manipulate easily (SESAME)

SESAME was designed to be a simple-to-use 3D design tool for 3D design. It is an ideal tool in the context of CoViD, as in any collaborative setting the computer skills of the participants will vary wildly.

The user interface of (the single-user version of) SESAME consists of a main 3D scene view and a menu panel on the right side. The menu offers a color/texture palette (top right), a 2D and 3D primitive shape selection palette (left part of menu panel), an undo button (at the bottom of the 3D palette), a navigation mode switch button, and a recycle bin (Fig. 1). The 2D palette consists of common 2D tools such as lines and arcs, and a freeform drawing tool. The 3D palette provides a tool to move objects in 3D and the instantiation of primitive 3D shapes, such as boxes, triangular prisms, spheres, and cylinders.

The user interface utilizes only a three-button mouse and a few modifier keys. All actions such as sketching a contour or moving objects can be accomplished with them. From a high level, the tool has two main modes: 2D and 3D, each of which is activated once any tool from the corresponding palette is selected. However, there are several operations that are available in both modes. An overview of the mouse function assignments is shown in Table 1.

Cloning provides a powerful way to create repetitive patterns in combination with the grouping technique

explained below. It is implemented as a (continuous) dragging action. That is, once a user drags the selected source object more than certain distance, a cloned object is instantiated. The cloned object will continue to move in the scene, until the user places it on the target position by releasing the mouse.

Additionally, SESAME provides a simple navigation interface so that users can assess the 3D structure of the scene rapidly. Camera rotation, pan, and zoom are accessible through middle mouse button dragging, shift-dragging, and scrolling, respectively.

2.1.1 Improved tools for 3D design

Beyond the basic actions listed above, SESAME provides also several advanced facilities for the rapid creation of 3D content.

With the 2D drawing tools the user can draw lines, arcs, or free-form curves onto any planar surface. As the user draws these in a 3D perspective view, SESAME displays a perspectively distorted circle during 2D drawing operations to help the user perceive the orientation of the current drawing plane. That is, in addition to visualizing the current line with the well-known rubber band technique, a circle is displayed with the origin at the start of the line and a radius proportional to the current length of the line (see Fig. 2).

Furthermore, multiple guides are displayed as well. These include guides parallel to the coordinate axes, as well as suggestions for lines that are perpendicular and parallel to other parts of the 2D drawing. The cursor snaps to these suggestions to aid the construction of common configurations. To avoid overloading the user, the SESAME displays only suggested completions that are close to the cursor.

Whenever the user adds any 2D shape, SESAME analyzes the current drawing and detects all closed contours created by the added shape. This also allows the user to subdivide existing closed contours by drawing other shapes over it. This effectively accommodates the creation of emergent shapes as these have been shown to facilitate creative thinking during the design process (Gross and Do 2000; Gross 2001). The user can then extrude any closed contour by clicking with the right mouse button inside the contour and dragging. For this, the height of the extrusion is proportional to the length of the drag, i.e. the top surface of the extrusion follows the cursor. A drag operation outward relative to existing 3D geometry will extrude a new volume (or resize an existing volume), while a drag inward will sculpt the volume, i.e. can be used to create holes. Figure 3 illustrates the technique.



Fig. 1 SESAME user interface

Table 1 Mouse and keyboard commands

Input command	Mode	Function	Realization
Left click	2D	Select 2D shape	On release
Left drag	3D	Select or add 3D object	On press
	2D	Draw 2D shape	
Shift + left drag	3D	Move 3D object	On dragging more than a threshold distance
	2D	Clone 2D shape	
Ctrl + left drag	3D	Clone 3D object	On dragging more than a threshold distance
	2D	Rotate 3D object	
Middle button	Both	Navigation	During dragging
Right drag toward outside of a volume	Both	Extrude a 2D contour or a face	During dragging
Right drag toward inside of a volume	Both	Subtract an extruded shape from a 2D contour or a face	Extrude during dragging, subtract after release
Shift + right drag	Both	Scale object	During dragging




Fig. 2 A circle is used to visualize the orientation of the current plane in SESAME. The user is currently drawing line *a*. Additionally, multiple colored guides are displayed based on other connected lines (*b* and *c*). The yellow line is perpendicular to *b*, the purple line is parallel to *c*, and the straight white lines are parallel to the coordinate axes

SESAME also provides a technique to let the user grab any object and slide it across the scene to the desired position to visually assess the impact of a change. One of the main motivations behind this is that this will greatly facilitate exploration. As almost all objects in the real world are attached to other objects, we based the design of our interaction technique on the idea that, unless special actions are taken, objects should *always* stay connected with another part of the scene and should not interpenetrate other objects. This conforms better to the way most people think about the real world. In contrast to other techniques, SESAME includes a novel technique that uses the entire area of the visual overlap of a foreground object with the (potentially complex) background, as this has been shown to work very well for 3D object manipulation (Oh and Stuerzlinger 2005). An evaluation showed that this new technique is significantly more efficient for

novice users and showed that our technique conforms very well to users' expectation about the position of objects relative to a scene.


Finally, SESAME also provides support for the selection of groups of objects. For this, the tool analyzes the scene and detects which objects are placed on top of another, as such objects will move with the base object in the real world. The implementation then simply selects objects on top of a base object whenever the base is selected (e.g. with clicking or rectangle selection). This turns out to greatly facilitate the manipulation of common scene configurations (Oh et al. 2006a). Lastly, a multi-click scheme, similar to character, word and paragraph selection in MS Word, cycles through an object with all objects it supports, the group consisting of all touching objects, and the object itself. Together, these selection techniques make restructuring and combinations of many 3D scenes easier.

2.1.2 Evaluation of SESAME

Traditional computer design tools are targeted toward the final stages of design. Senior designers have commented, “recent designs are very beautiful due to the (use of a computer) design tool, but they are very poor, far from the design ideal” (Tano et al. 2003). This means that while all details are perfect, the users of current design tools often do not adequately explore the space of all possible design solutions.

SESAME was previously evaluated to analyze its suitability as a design tool. A comparison with sketching with pen and paper revealed that SESAME yields the same level of design quality as judged by an expert designer (Oh et al. 2005). Another comparison with users of 3D Studio Max demonstrated that even first-time users of SESAME could generate meaningful

Fig. 3 Sketching interface. **a**, **b** Dragging outward from a closed contour creates a volume. **c**, **d** Drawing a contour on a face and dragging inward sculpts the volume. **d**, **e** The user creates a new shape by dragging outward. **e**, **f** The user can also “stretch” an object by dragging the face directly. (Red lines added to visualize drag operations)



designs faster with this tool than traditional CAD systems (Oh et al. 2006b). Also, studies with naïve users (i.e. people who had never used 3D design systems before) showed that most people quickly understand how SESAME works and are rapidly able to use it to create 3D designs.

SESAME, as described, is capable of creating mainly prismatic objects, i.e. everything that can be modeled by extrusions. This limits the application domain to engineering domains, interior design, as well as architecture. Current work on SESAME aims to add support for rotationally symmetric objects as well as freeform modeling tools.

2.2 Multi-user laser table interface (MULTI)

In general, the stakeholders in a design project make important high-level decisions in a meeting. However, for effective collaboration, the participants in such a meeting need to be able to see a visualization of the current design proposal(s). If this visualization is interactive, the participants will in general be able to reach better decisions, as they can explore the design space more completely.

However, 3D design activities frequently require large displays, as they make it simpler to assess the impact of a proposed design. Another factor is that screen space is often limited in many situations, especially when multiple proposals are evaluated simultaneously. Furthermore, in collaborative meetings, there is always secondary information that is critical to the task, but not represented on the main visual display. Often this information is available in printouts, but this is clearly not a very interactive medium. Another alternative is various kinds of other mobile devices, but all of these are essentially single-user devices and are

often poor platforms for collaborative activities. As discussed above, most meetings take place around tables. However, Ryall et al. 2004 stated that for groups with more than three persons a single tabletop is not sufficient and that larger groups may well need vertical displays for shared information. Hence, it is desirable to have an infrastructure that supports both an interactive tabletop as well as several additional large interactive display surfaces.

The MULTI infrastructure was designed for collaborative 3D design. It consists of an interactive table and three interactive walls that are positioned in a semi-circular arrangement around one of the short sides of the table. A single 3 GHz Windows XP computer, with three graphics cards with two outputs each, drives all five projectors. The infrastructure includes also a standard 802.11 wireless network and several Tablet PC's (not shown). An image of MULTI is shown in Fig. 4.

MULTI is designed for groups of five to seven active participants and the size of the table was chosen so that five people can comfortably sit around it. To give users full access to both long sides and one of the short sides of the table, we had to use two projectors for the table, each of which covers approximately half of the table. The decision to add several walls was motivated by several factors: the desire to be able to work on large-scale designs, the desire to support even more users as a “passive” audience, and the ability to use e.g. only the walls alone for large-scale visualization. Furthermore, the table has a stable ledge around the interactive surface in the middle, which provides space for laptops, paper, and other work artifacts (and is even stable enough to sit on).

To avoid issues with shadows cast by user's hands etc., all surfaces are back-projection screens and the



Fig. 4 An overview of the screen configuration of MULTI. All projectors and cameras are located behind the screens. The image on the table is created by two projectors (partially located underneath the middle wall screen)

projectors and cameras are placed behind the screens (both for the table and the walls).

Each of the three back-projection screens has a 60" diagonal and is illuminated by a NEC WT600 projector, which have an extremely short throw distance. This helps to reduce the space usage of the configuration. The tabletop has also a 60" diagonal and two WT600 projectors illuminate it from underneath (the front and back half of the table, respectively). This "split-screen" design gives users free access to the two long sides and one short side of the table. Careful positioning of the projectors guarantees an almost seamless image. All projection screens are acrylic with a 0.7 gain. A pane of tempered glass thick enough to carry the weight of an adult supports the tabletop screen and a thin protective layer on top prevents scratches. Finally, there is a 15 cm wide metal ledge around the three accessible edges of the tabletop.

The main interaction devices of MULTI are several computer-controlled laser styli. The laser spots observable at any intersection of the laser beam with the projection surfaces are detected by a set of cameras. Each stylus works both as a pen (as in pen-based computing) as well as a remote pointing device. The button on a stylus is configured to work as the left mouse button, which makes it very natural to interact with standard GUI applications. Each laser stylus is wired to a control box that allows the computer to control their laser diode and the status of the button on the stylus. The laser diodes are driven with a time-multiplexing scheme that allows the infrastructure to identify each laser stylus (Pavlovych and Stuerzlinger 2004). In the MULTI system, we use a "scaled version"

of this time-multiplexing scheme that allows simultaneous tracking of seven laser styli on all four surfaces. All laser spots are tracked with Naturalpoint cameras, which feature partial hardware support for the detection of bright spots. We removed the IR filter from each camera to enable tracking of the laser spots via the standard CMOS camera chip. Each camera has a 57° lens, but due to sightline restrictions MULTI uses six cameras underneath the table, four behind the middle wall, and one behind each sidewall. As the time multiplexing of the laser diodes is synchronized with the refresh rate of the cameras (120 Hz), the blinking is (usually) not visible to the human eye.

The software tracking the laser styli can either send all stylus events to multi-user aware applications over a socket interface or alternatively simulate standard mouse events for the operating system. This second alternative is somewhat limited as the operating system silently assumes a single input device and displays only one cursor. This facility is very useful, however, as it enables the users to interact with the standard GUI environment of the operating system and all installed applications. Further details about the construction, the hardware, and software for MULTI will appear in a forthcoming publication.

2.3 CoViD: merging SESAME and MULTI

The three main modifications necessary to enable the functionalities of SESAME to run on MULTI were the support for multiple display windows, adaptations to support laser styli as input devices, and the support for multiple simultaneous users. Together, this forms the CoViD system.

First, the displays of MULTI were associated with various views of the 3D design. The tabletop surface is best used for a view of the design from the top, akin to e.g. the view afforded by a map of a city. The vertical surfaces are better suited for a "side" view of the 3D model and we typically use two or all three wall displays together to provide a large-scale perspective view of the design (see e.g. Fig. 5). As additional information is critical in many design meetings, we allow the users to display and interact with additional information (e.g. a browser window or a spreadsheet) on one (or more) of the wall screens a browser window. These windows are typically placed on one of the two "side" walls. Furthermore, several tablet PC's are available in the system as another means to access auxiliary information or to allow people to transition between group and individual work. Connectivity is provided via a wireless network and two USB hubs, which are mounted in convenient locations under the table.




Fig. 5 The CoViD system in action

Technically speaking, it was fairly easy to adapt SESAME to support multiple views. We only had to change the rendering loop to support multiple open windows. However, there were subtle issues related to limitations of graphics card drivers, which prevented us from creating a single window across that spans all three wall displays. We currently handle this case simply by opening one window on two displays and another on the third.

The second issue, supporting a laser stylus as input device, was done by modifying the central event loop of SESAME to check for updates from the laser stylus tracking software of MULTI, which sends all information about stylus movements and button presses over a socket. Then, each stylus event is handled by the standard event handling code in SESAME, similar to how mouse events are handled.

However, a limitation of the current MULTI laser styli is that they have only a single button, while SESAME assumes a three-button input device. The middle mouse button is used for navigation, while the right mouse button is used for extrusion, sculpting, and resizing. On the other hand, it doesn't make much sense to allow all users to change the camera for the shared tabletop display as this would allow one user to change the view while another is manipulating content. Hence, we decide to support navigation only via a single wireless mouse attached to MULTI, which naturally enforces a turn-taking protocol. The functionality of the right mouse button in SESAME is important for the design process. Hence, we added a special icon to the tool palette of SESAME, which allows users to toggle “right-button” mode for their stylus (and also added appropriate visual feedback close to the cursor position).

To support multiple operations by multiple users at the same time, we enhanced SESAME to keep track of separate states for each user. Also, the software needs to simulate a separate cursor for each participant. Whenever a user selects one of the tool palette entries, his/her cursor is changed to reflect their current mode. Furthermore, the undo functionality of SESAME was adapted by allowing people to activate undo in any display.

Finally, we had to perform more extensive modifications to SESAME to enable the creation of new content in parallel by multiple users. This involved keeping track of the currently active “drawing plane” on a per user basis and modifying that information with every 2D interaction of a user. We also adapted the 2D drawing facilities of SESAME so that all guides are only visible in the current window—i.e. do not show up in other windows. This avoids obscuring other views of the 3D model with these temporary visualizations.

3 Evaluation of CoViD against guidelines

This section presents an evaluation of the CoViD system relative to the guidelines presented in the first part of the paper. This evaluation was performed by one of the authors, a human–computer interaction expert. As most points focus on one of the two main parts of CoViD (SESAME and MULTI), we refer these parts whenever appropriate.

- *Non-intrusive interface:* To minimize the cognitive load of the user SESAME relies only a minimum set of modes: a 2D mode for drawing contours and a 3D mode for manipulating solids and the various tools associated with each mode. For 2D drawing activities, the system provides a rich set of support techniques—e.g. suggestions, automatic segmentation of freehand drawings, and recognition of closed structures. In 3D manipulation mode, the system provides a natural and efficient manipulation techniques based on physical properties, such as gravity and collisions. To avoid the cognitive overhead associated with wire-frame visualization and/or orthogonal perspective, SESAME allows the user to perform all manipulations in a perspective view.
- *Easy creation:* As solid modeling is more appropriate for early design phases compared to other approaches such as polygonal modeling (Deisinger et al. 2000), SESAME supports solid modeling directly, via the metaphor of the extrusion of 2D contours, as well as sculpting and direct manipulation of 3D objects.

- *Easy combination and restructuring:* SESAME provides object manipulation schemes that use a single viewpoint and match the users' expectations about the most probable 3D object motion. Similarly, group manipulation techniques behave in a way that is consistent with the real world. Both of these techniques facilitate experimentation with the structure of the scene, as objects behave like most naïve users expect them to do.
- *Tolerance to ambiguity and incompleteness:* Sketching naturally supports ambiguity and incompleteness, which makes it a great tool for the initial phases of design. SESAME includes support for three of the most important features of sketching. First, any freehand sketch is automatically beautified by segmenting it into primitive 2D shapes. As user input usually contains some amount of jitter, often also caused by digitization, a certain degree of beautification is usually a positive thing. In SESAME, the level of detail for the segmentation varies depending on the view distance to preserve important details while removing unwanted noise. We found that this addresses the trade-off between ease-of-use and precision fairly well. Also, the parts of segmented stroke can be manipulated later, i.e. if the user wants to refine a design, which facilitates the transition to more precise designs. Second, SESAME supports the coexistence of 2D drawings and 3D shapes, as it facilitates the “exchange” of ideas across dimensionalities. Finally, SESAME renders the scene as flat-shaded objects with thick outlines overlaid on top of them. This is a compromise between the “perfect” look of standard computer graphics and the “imprecise” look of a pen-and-paper sketch. While it would be desirable to provide a more “sketchy” look, the problem is that there are no *efficient* and *general* interaction techniques that work with imprecise 3D information.
- *Range of levels of abstraction:* Snapping provides a convenient bridge between the inaccuracy of a users' input and the accuracy required by computer system. With snapping, designers can quickly generate a rough scene configuration and the system can display it accurately (enough). In SESAME, the level of detail is proportional to the viewing distance in many interaction techniques. This is based on the observation that designers generally work on an overall idea by viewing the scene from a farther distance and work on detail by zooming into the part of interest. Furthermore, the group manipulation techniques of SESAME facilitate fluid transitions between work at a large scale and detail modification.
- *Ability to edit various forms of information:* Currently, SESAME does not address this in a significant way.
- *Supporting evaluation (simulation):* We support this is via the export of scene geometry to any simulation package (e.g. photorealistic rendering software or stress analysis) and then running that simulation. In the future, we may integrate this functionality better.
- *Natural interpersonal interaction:* As all participants are sitting (or standing) around the MULTI table, this is similar to standard meetings around a conference table. Hence, participants can naturally interact face-to-face with other participants in the system.
- *Transitions between activities:* MULTI allows its users to interact with the standard GUI desktop, which enables them to use many different applications and hence transition seamlessly between activities such as design, performing a WWW search, interacting with spread-sheets or other documents, etc. As the wall surfaces are interactive as well, it is easy to transition to sessions, where one person interacts with the system as if it were a large blackboard. Furthermore, people can (and usually do) place paper documents and other work artifacts on the rim of the table, which provides even more possibilities for transitions between various activities.
- *Transitions between personal and group work:* The table surface of MULTI is large enough that a person cannot simply reach completely across it (without standing up and leaning across). This effectively gives each user their own personal space in an unambiguous manner. Furthermore, all work files are placed in subfolders of the GUI desktop. As this desktop folder of the system is shared via the wireless network, this makes it easy for participants to transfer (part of) the data to an external portable device (e.g. a Tablet PC) for individual work and back to the desktop to share it again.
- *Transition between tabletop collaboration and external work:* Two USB hubs have been mounted in easy to access locations under the table. This facilitates transfer of data from USB drives to and from MULTI. Furthermore, the wireless network as well as the standard network connection of the system provide for even more opportunities to transfer data to and from MULTI.
- *The use of physical objects:* The wide metal ledge around the table is stable enough for a user to sit on and users can place laptops, various printouts, notebooks, and other personal items on it. While

people may be tempted to put a lot of stuff around the ledge, this may also block their access to the table, which balances things nicely. The interactive tabletop itself is stable enough to support considerable weight (due to the pane of tempered glass behind it). This allows people to place objects onto the interactive part, but such objects block the projection. Currently, MULTI does not support a tangible user interface, but we are working on adding this functionality.

- *Accessing shared physical and digital objects:* To access a shared digital object the user just needs to point at it with his/her laser stylus. In computer-based design applications, the orientation of artifacts on the table surface is less of a concern, as the natural perspective for content on the table surface is a top down view onto the content. This makes it relatively easy to access objects from any direction. The wall surfaces of the system provide a side view of the 3D environment, but these surfaces have anyways a natural “up” orientation. All physical/tangible objects on the table surface can easily be manipulated by reaching for them. Finally, all interactive surfaces are back-projected in MULTI, which makes it impossible for a user to cast shadows onto content (beyond normal blocking of shared content with a hand).
- *Flexible user arrangements:* Due to the size and the physical configuration of the table (only one short side is blocked), users can sit in many configurations around it. As the table is on heavy-duty rollers, it can even be rolled away completely if the users only want to only use the wall surfaces. As discussed in the previous point, in the context of design applications, the natural view for the table is top-down, which leaves a lot of flexibility for the users to arrange themselves around the table.

- *Simultaneous user interactions:* Due to the time multiplexing scheme, several people can simultaneously interact with MULTI. Due to the restrictions of the Windows operating system (it only supports one cursor), standard GUI applications do not support multiple users seamlessly. However, applications that include support for multiple input devices can use the full functionality of MULTI.


4 Pilot user study of CoViD

We conducted two pilot experiments to test the basic features of CoViD in single user mode in comparison with collaborative mode. Nine participants were recruited from the pool of computer science undergraduate and graduate students (two female, seven male) at the local university. They were organized in three groups of 3. Each participant was asked to perform two tasks individually as well as collaboratively in the group. Before the experiment, we explained how the system worked and demonstrated the use of the laser styli as well as the relevant operations of SESAME to each participant for about 10 min. At the end the experiment, users were given a questionnaire to rate preferences between working in groups versus working individually and which task they enjoyed the most. Statistics was analyzed via ANOVA.

4.1 Scene assembly task

We first evaluated the effectiveness of the system to assemble a simple object with the 3D movement technique of SESAME. The participants were asked to assemble a chair from parts first individually, as shown in Fig. 6.

Fig. 6 Scene assembly task. **a** Initial state. **b** Target scene



4.1.1 Results

When working individually, average completion time for the task was 4 min and 23 s. When working collaboratively the average completion time was 5 min 40 s. Detailed timings are shown in Table 2.

None of the differences between the individual and group performances are statistically significant.

4.2 Design task

In the second task we asked users to create a simple 3D design. The task was to create three different 2D shapes using the rectangle, freeform and circle tools, to turn them into 3D structures via extrusion and then create “chimneys” on top of them. The condition was that the chimneys had to be created using tools different from the ones that were used to create their corresponding structures, as illustrated by Fig. 7.

4.2.1 Results

When working individually, average completion time for the task was 4 min and 5 s. When working collaboratively, average time was 4 min and 54 s. Detailed timings are shown in Table 3.

Again, none of the differences between the individual and group performances were statistically significant.

Table 2 Scene assembly times by user (in seconds)

Individual	240	191	199	440	314	480	139	231	110
Average individual	210			441			160		
Group	300			420			288		




Fig. 7 Sample configuration for the design task

Table 3 Design task times by user (in seconds)

Individual	420	300	179	434	153	480	148	163	107
Average Individual	300			355			139		
Group	300			413			171		

4.3 Discussion

Due to the small sample size (only three groups) none of the results of the statistical analysis were significant, which limits our ability to draw conclusions from the data. One factor that may have contributed to this is that the tasks were relatively simple and the potential benefits of collaboration hence negligible. However, we can still state that participants were able to complete the tasks reasonably quickly.

One interesting behavior we observed was that two groups worked mainly in the perspective view and used the tabletop only when users wanted to work in parallel. This may have been based on the relatively larger size of the scene on the wall screens. The lack of a second button on the laser styli was also an issue that led to unnecessary errors.

The analysis of the questionnaires revealed that on average participants rated the four different conditions (individual scene assembly, individual design, collaborative scene assembly, and collaborative design) about equally in terms of how much they enjoyed the different conditions.

As for the collaborative assembly task, users commented that assembling parts with others made it more confusing because complained that other users would often interfere and “ruin” their work. For collaborative design, users commented that working with other people made the task more entertaining and that it was more fun to create new designs together rather than assembling parts. This hints at the potential of CoViD to support creativity in design sessions. Some users reported problems with the laser styli in the palette area on the table. This seems to have been caused by calibration problems in that region of the tabletop. The small seam (less than one pixel) caused by the overlap between the images of the two table projectors was not recognized as a noteworthy issue.

The evaluation of CoViD revealed several issues that need to be addressed, before a more comprehensive user study can be performed. Most importantly, we need to redesign the laser styli to be wireless and to include a second button. Furthermore, there were calibration problems in certain areas of the screens and some synchronization problems between the laser

diode circuit and the camera systems, mainly due to some bugs in the camera driver. We are currently working on fixing these issues.

5 Conclusions and future work

This paper presented CoViD, a novel system for collaborative 3D design. We first introduced a set of guidelines for collaborative design systems. Then we introduced the new system and presented details about its components and how they were adapted to form a collaborative 3D design system. Subsequently, we evaluated the design of the system against the guidelines and reported results of a pilot study, in which users generally liked the ability to collaboratively design in 3D.

We are currently working on addressing the issues that were identified during the pilot study. Once that is done, we intend to perform a more comprehensive evaluation by having a group of architectural students design building(s) for a free lot in a collaborative setting in collaboration with a professor in architecture. Furthermore, we also intend to analyze how the collaborative aspect of the CoViD system affects creativity. Another area of future work is the extension of SESAME to support additional modeling primitives (such as rotationally symmetric and freeform objects). Finally, we are working on a tangible user interface for MULTI and plan to extend SESAME to allow for tangible interaction.

Acknowledgements Thanks to D. Phillips, D. Dadgari, and A. Vorozcovs for help with camera calibration and programming the laser spot detector software, to the York Centre of Vision Research, and to NSERC for funding.

References

- Deisinger J, Blach R, Wesche G, Breining R, Simon A (2000) Towards immersive modeling - challenges and recommendations: a workshop analyzing the needs of designers, Virtual environments, Springer, Berlin Heidelberg New York, pp 145–156
- Druin A, Stewart J, Proft D, Bederson B, Hollan J (1997) KidPad: a design collaboration between children, technologists, and educators. *Conf ACM CHI* 97:463–470
- Gross M, Do E (2000) Drawing on the back of envelope: a framework for interacting with application programs by freehand drawing. *Comput Graph* 24:835–849
- Gross M (2001) Emergence in a recognition based drawing interface. *Conf Vis Spatial Reason*, 51–65
- Igarashi T, Matsuoka S, Tanaka H (1999) Teddy: a sketching interface for 3D freeform design. *Conf ACM SIGGRAPH'99*, 409–416
- Oh J-Y (2005) Desktop 3D conceptual design systems. PhD Thesis, York University
- Oh J-Y, Stuerzlinger W (2005) Moving objects with 2D input devices in CAD systems and desktop virtual environments. *Conf Graphics Interface*, 141–149
- Oh J-Y, Stuerzlinger W, Danahy J (2005) Comparing SESAME and sketching for conceptual 3D design. *Workshop EG Sketch Based Interfaces Modeling*, 81–88
- Oh J-Y, Stuerzlinger W, Dadgari D (2006a) Group selection techniques for efficient 3D modeling. *IEEE Symp 3D User Interfaces*, 95–102
- Oh J-Y, Stuerzlinger W, Danahy J (2006b) SESAME: towards better 3D conceptual design systems. *Conf ACM Dis*, 80–89
- Pavlovych A, Stuerzlinger W (2004) Laser pointers as interaction devices for collaborative pervasive computing. *Adv Pervasive Comput*, 315–320. ISBN 385403176-9
- Ryall K, Forlines C, Shen C, Morris M (2004) Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. *Conf ACM CSCW*, 284–293
- Scott S, Grant K, Mandryk R (2003) System guidelines for co-located, collaborative work on a tabletop display. *Conf ECSCW*, Springer, Berlin Heidelberg New York, pp 159–178. <http://www.springer.com/east/home/generic/search/results?SGWID=5-40109-22-33658793-0>
- Shu L, Flowers W (1992) Groupware experiences in three-dimensional computer-aided design. *CSCW*, 179–186
- Stewart J, Bederson B, Druin A (1999) A model for co-present collaboration. *Conf ACM CHI*, 286–293
- Tano S, Kodera T, Nakashima T, Kawano I, Nakanishi K, Hamagishi G, Inoue M, Watanabe A, Okamoto T, Kawagoe K, Kaneko K, Hotta T, Tatsuoka M (2003) Godzilla: seamless 2D and 3D sketch environment for reflective and creative design work. *IFIP INTERACT'03*, 311–318
- Tse E, Greenberg S (2004) Rapidly prototyping single display groupware through the SDGToolkit. *Australasian User Interface Conf*, 101–110
- Zelevnik R, Herndon K, Hughes J (1996) SKETCH: an interface for sketching 3D scenes. *Conf ACM SIGGRAPH'96*, 163–170